Experiences with Operations and Autonomy of the Mars Pathfinder Microrover

Andrew H. Mishkin
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
818-354-0986
Andrew H. Mishkin@jpl.nasa.gov

Jack C. Morrison
Telos Systems Group
320 N. Halstead St. Suite 260
Pasadena, CA 91107
818-354-7782
Jack.C.Morrison@jpl.nasa.gov

Tam T. Nguyen
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
818-354-7296
Tam.T.Nguyen@jpl.nasa.gov

Henry W. Stone
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
818-354-9051
Henry.W.Stone@jpl.nasa.gov

Brian K. Cooper
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
818-354-6298
Brian.K.Cooper@jpl.nasa.gov

Brian H. Wilcox
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
818-354-4625
Brian.H.Wilcox@jpl.nasa.gov

Abstract— The Microrover Flight Experiment (MFEX) is a NASA OACT (Office of Advanced Concepts and Technology) flight experiment which, integrated with the Mars Pathfinder (MPF) lander and spacecraft system, landed on Mars on July 4, 1997. In the succeeding 30 sols (1 sol = 1 Martian day), the Sojourner microrover accomplished all of its primary and extended mission objectives. After completion of the originally planned extended mission, MFEX continued to conduct a series of technology experiments, deploy its alpha proton x-ray spectrometer (APXS) on rocks and soil, and image both terrain features and the lander.

This mission was conducted under the constraints of a once-per-sol opportunity for command and telemetry transmissions between the lander and earth operators. As such, the MFEX rover was required to carry out its mission, including terrain navigation and contingency response, under supervised autonomous control. For example, goal locations were specified daily by human operators; the rover then safely traversed to these locations. During traverses, the rover autonomously detected and avoided rock, slope, and drop-off hazards, changing its path as needed before turning back towards its goal. This capability to operate in an unmodeled environment, choosing actions in response to sensor input to accomplish requested objectives, is unique among robotic space missions to date.

This paper describes the techniques implemented on MFEX for operations and autonomous control; the performance of this vehicle on Mars is also discussed.

TABLE OF CONTENTS

1. INTRODUCTION

On July 4, 1997, the Pathfinder spacecraft successfully landed in the Ares Vallis region of Mars. On sol 2, the second Martian day after landing, after overcoming an initial problem with poor lander-rover communication, the Sojourner rover rolled down the rear lander ramp onto the surface. In order to carry out its mission, traversing to sites of scientific interest, the rover design incorporated autonomous capabilities not applied in previous planetary exploration missions. The mission called for rover traverses to be performed nearly every sol, requiring Earth-based operators to build sequences in hours, rather than the days necessary during Viking, the previous U.S.-conducted landed mission.

This paper provides an overview of the rover mission objectives, and the implementation of the rover intended to meet those objectives. We discuss how day-to-day operations were conducted to carry out the mission. To provide context for the autonomous capabilities of the

Sojourner rover, we describe the steps performed by the human operators, defining where autonomy was used and where it was not. Then we evaluate our experience operating Sojourner during the first 70 sols on the surface of Mars.

2. ROVER MISSION OBJECTIVES

As of this writing, all rover mission objectives have been met or exceeded. These objectives included:

Cost and Schedule

The MFEX budget, including design, development, implementation, and operations, was $25M. All critical delivery deadlines for integration of the rover elements with the Pathfinder mission were met, culminating in the launch on December 4, 1996.

Mass

The mass allocation for the rover and its lander-mounted support equipment (tie-downs, rails, ramps, and UHF radio link) was 16 kilograms, while the actual combined mass of all rover elements is 15.2 kilograms. The mass of the rover itself is 10.5 kilograms.

Rover Impact on Pathfinder Project

The interfaces between rover and lander were simplified as much as possible to reduce dependencies between the two development efforts. For example, there was no electrical interface between the rover and the lander. To wake up the rover during pre-launch and cruise mission phases, a reed relay switch in the rover was activated by a magnetic coil mounted on the lander petal; activation of the switch allowed the rover's own batteries to power its bus. For telemetry processing, the rover transmits already formatted packets to the lander, which then processes them in the same manner as packets generated by the lander itself.

Survivability

The rover had to survive the launch, cruise, landing, and Mars surface environments to which it would be subjected.

Surface Operations Objectives

Rover mission success was defined primarily by the accomplishment of the surface operations objectives. One complete set of technology experiments, including soil mechanics, material adherence, and wheel abrasion, together with one APXS rock data collection, and an image of the lander to assess its post-landing condition, have been defined to constitute 90% of mission success. The remaining 10% is achieved by completing additional sets of technology experiments, APXS data collections, and imaging activities.

Operating Range

The mission plan called for the rover to operate primarily within 10 meters of the lander; this is considered the effective limit of usefulness of the lander stereo images for directing the rover and identifying sites of scientific interest. If desirable destinations for the rover are identified further from the lander (in particular during the extended mission), then the rover may be commanded to travel as far as the lander's horizon. The rover's design allows it to drive several hundred meters from the lander before passing out of communications range. The software design enables it to respond to communications loss in one of two specified ways: 1) stop and back up to re-establish communication, or 2) continue executing its sequence, which will bring the rover back into communications before it completes. (If human error results in a sequence that ends with the rover outside of communications range, an onboard contingency sequence will be triggered, causing the rover to drive toward the origin of its lander-centered coordinate frame.) While this long distance driving is feasible because the rover's architecture, the rover's hardware has been qualified to ensure not less than 100 meters of traverse on the Martian surface.

Lifetime

The rover's prime mission has been designed to allow the rover to accomplish its surface operations objectives in the first seven sols of operations. In addition, no element of the rover's design should preclude its operation for a full 30 sol extended mission, during which greater risks may be taken. The only exhaustible resource (other than normal wear) is the non-rechargeable battery; the rover is capable of performing its entire mission, with the exception of night time APXS data collection, even if the batteries are unavailable after landing.

3. ROVER DESCRIPTION

The rover design and implementation have been described previously in references [1], [2], [3], [4]. The MFEX rover is a six-wheeled robotic vehicle that is 68 cm long by 48 cm wide, standing 28 cm high when fully deployed (see Figures 1 and 2). It has 13 cm diameter wheels. After deployment, the vehicle has a ground clearance of 17 cm.

The mobility subsystem consists of a 6-wheel drive, 4-wheel steerable, rocker-bogie mobility chassis. This configuration allows the vehicle to surmount obstacles 1.5 wheel diameters in height. The rover's speed is approximately 0.4 meter/minute in nominal terrain.

The rover computer possesses a single CPU, an Intel 80C85 operating at 2 Mhz, processing 100 KIPS (thousand instructions per second). Four types of memory are incorporated into the electronics boards: 16 Kbyte radiation hardened PROM, 64 Kbyte radiation hard RAM, 176 Kbyte

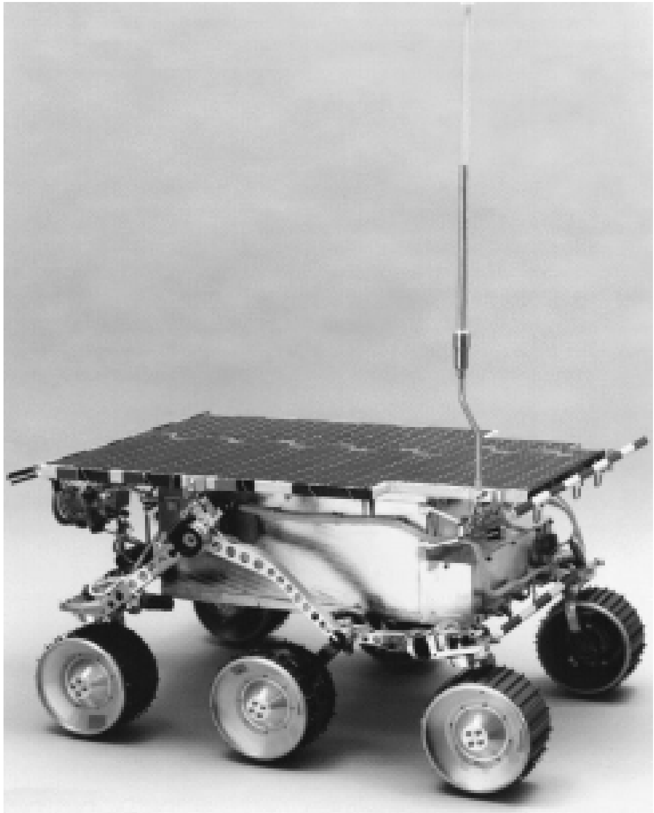

Figure 1. MFEX Rover

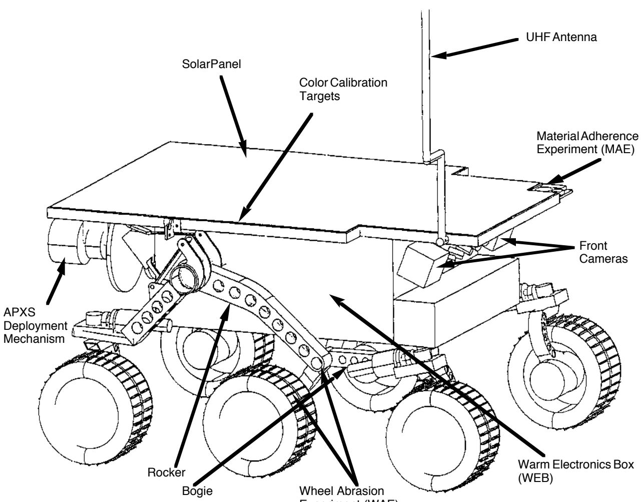

Figure 2. Major Elements of the MFEX Rover

bulk EEPROM, and 512 Kbyte RAM, addressable in 16 Kbyte pages. Two double-sided printed wiring boards implement all processing and power conditioning/ distribution.

Black-and-white stereo cameras on the front of the vehicle support hazard detection and science/operations imaging. The image size for each camera is 768 x 484 pixels, with a 4 mm, wide-angle lens providing a field of view of 127° x 94° and a resolution of 3 mradians/pixel. Each camera is mounted about 25cm above the terrain surface. A single color rear camera is used primarily for APXS target documentation.

Power for operations is supplied by a Gallium Arsenide (GaAs) solar panel that provides 15W peak; primary (non-nonrechargeable) batteries enable night operations of the APXS instrument, provide extra power for mobility if needed in difficult terrain, and serve as a backup in the event of solar panel failure on landing.

A pop-up antenna is located at the edge of the solar panel. Communications between lander and rover is via UHF radio modems with a raw data rate of 9600 bits per second. Overhead for the link results in an effective data transfer rate of approximately 2 Kbits per second. Maximum separation between rover and lander for communications is 500 meters.

To protect the onboard electronics and batteries from the temperature extremes of the Martian surface, they are housed in the Warm Electronics Box (WEB). The WEB is of face sheet and spar construction, with solid silica aerogel insulation. The WEB is heated by a combination of waste heat produced by the electronics during normal daytime operations, heat generated by 3 radioisotope heater units (RHU's) providing nearly 3 watts, and by channeling excess available solar power to internal heaters. The WEB cools overnight, then begins to warm again in the early morning when the rover powers on. This strategy maintains the electronics in the ±40°C flight allowable temperature range for the duration of the mission.

Onboard experiments include: the Material Adherence Experiment (MAE) supplied by the Lewis Research Center (LeRC), which measures solar panel output; the LeRC Wheel Abrasion Experiment (WAE) on the right center wheel and bogie, which investigates the abrasiveness of the Martian soil; and the APXS, which determines the elemental composition of rocks and soil (provided by the University of Chicago and the Max Planck Institute in Mainz, Germany). Other experiments rely on the rover's imaging, engineering, and navigation sensors to generate data necessary for analysis.

4. SURFACE OPERATIONS SCENARIO

The rover operations team prepares one command sequence per sol (one Martian day). The design of each sequence is based on a combination of 1) the rover state assessment provided by the Rover Engineering Analysis Team, 2) the science and technology experiment requests from the Experiment Operations Team, and 3) the feasibility of the requested operations given the trafficability of the Martian terrain and the safety of the vehicle. The uplink team designs a sequence to fulfill as many of the science and technology requests as possible while maintaining the health of the rover.

On a given sol, there is usually only one opportunity to uplink rover and lander command sequences. This opportunity corresponds to early- to mid-morning of the Martian day.

Telemetry is commonly downlinked during three periods per sol. The first downlink occurs just prior to the morning uplink. The mission operations team has a short time to review the telemetry to determine whether any contingencies have occurred during the Martian night that would preclude uplinking the nominal sequence. The second downlink is around mid-day, before the rover has completed its primary operations for the sol. The final downlink takes place mid- to late-afternoon on Mars, and provides the primary telemetry necessary to plan the rover's next sol's activities.

The rover nominally operates autonomously for one sol (>24 hours) until receipt of the next command sequence. During a typical sol, the rover will perform a subset of the following operations: complete an APXS data collection that was carried out during the prior night; capture a rear color image of the APXS site; traverse to an appropriate site and perform a series of soil mechanics experiments, including several subframe images of soil mounds and depressions created by running individual wheel motors; perform a WAE experiment and several MAE experiments; traverse to a designated rock or soil location; place the APXS sensor head; capture end-of-day operations images with its forward cameras; begin APXS data collection; and shut down for the night. APXS data collection usually occurs overnight while the rover is shutdown.

Once the rover has completed its traverse activities for the sol (usually by 2:00pm Mars local time), the IMP (Imager for Mars Pathfinder) camera on the lander captures one or more stereo images of the rover at its end-of-day location. These images are required for operations; they are therefore given a high telemetry priority to ensure that they are downlinked during the current sol.

Downlinked images and rover telemetry are used by the rover team to assess the rover's state and to plan the next sol's activities. In order to prevent the accumulation of dead reckoning error from sol to sol, the uplink team uses the lander images of the rover to localize and update the rover's location on the Martian surface. These images of the rover are merged with the stereo "monster pan" of the terrain which was built up over the first few sols of operation. Rover destinations are then designated in the stereo display of the Rover Control Workstation and integrated with the rest of the rover command sequence.

5. COMMAND SEQUENCE GENERATION

Typically, one command sequence defines the rover activities for one sol (including both day and night operations), plus "runout" commands in the event the next sequence is delayed. Traverse commands are only a small fraction of most command sequences. Other commands (or sets of commands) control the following functions: update of rover position and orientation; imaging; passing of commands to the APXS instrument; collection of soil mechanics experiment data, MAE data, and WAE data; deployment of the APXS sensor head; parameter settings for "housekeeping" functions such as heating times and self-diagnostic health check rates; rover shutdowns with appropriate wakeup times; error masking and clearing; contact sensor masking; low-level activation of devices and actuators (if needed); telemetry buffering options; enabling of battery usage for various devices.

Unlike previous planetary exploration missions, we needed to generate new command sequences within a few hours based on telemetry downlinked daily. Until we receive end-of-day images and engineering data that shows us where the rover has actually gone, and what its current state is, we cannot decide where the rover should go next, or what it should do in the process. To facilitate this rapid command turnaround, we have generated a set of command sequence macros for activities that the rover performs on a repetitive basis. These macros encompass experiment operations, APXS site imaging, final approach to rocks of different sizes, night time APXS data readouts, and common groupings of parameter settings.

Human operators design rover command sequences at the Rover Control Workstation (RCW). The operator can "fly" a 3-D rover icon through the stereoscopic display of the Martian terrain. By inspecting the stereo scene, as well as placing the rover icon in various positions within the scene, the operator can assess the trafficability of the terrain. By placing the icon in the appropriate position and orientation directly over the stereo image of the actual rover on the surface, we automatically compute the rover's location and heading; when this information is uplinked to the rover, accumulated dead reckoning error is corrected. The rover driver specifies the rover's destinations by designating a series of waypoints in the scene, generating waypoint traverse commands. Other types of commands are inserted into the sequence using a customized graphical user interface.

Since downlink communications data is a scarce resource, the telemetry volume produced by a given sequence must be estimated during sequence review to ensure that the rover does not exceed its allocation for the sol.

The RCW generates rover command sequence files in a format compatible with the Mars Pathfinder project uplink tools. These files are delivered to the project flight engineers, who insert them into the sequence load for the spacecraft. Once uplinked to the lander, the rover sequence is activated, placing it in the lander's rover buffer, to be provided to the rover the next time the rover requests a commandsequence.

6. AUTONOMOUS ROVER CAPABILITIES

The rover has autonomous capabilities in the areas of terrain traverse, contingency response, and resource management.

The rover can autonomously navigate through the natural terrain of the Martian surface to locations specified by Earth-based operators. The time-delay (10 minutes one way on July 4) intrinsic to communications between Earth and Mars precludes real-time human operator control of the rover. Continuous communications between Earth and Mars is not feasible in any case, because of Deep Space Network (DSN) contention issues. In addition, Earth-based operators viewing lander-based images of the scene may not be able to discern all hazards to the rover. Therefore, the rover must be able to respond to sensor input (from accelerometers, rate sensor, encoders, articulation sensors, and the proximity hazard detector) on its own in real time in order to reach sites of interest in reasonable time, as well as to protect itself from attempting hazardous, potentially mission-ending maneuvers.

Communications to the rover can potentially fail across two links: the lander receiver may fail, preventing rover sequences from being uploaded; and the local UHF link between the lander and rover may become inoperative. In the first case, the rover operates normally, because the lander will periodically activate pre-loadedbackup rover sequences. In the second case, the rover will eventually trigger its own on-board contingency sequence, and perform a generic mission, transmitting its telemetry without handshaking in the hope that the lander is still listening. (The rover's actions in these two situations is described more fully below in "Contingency Scenario Responses".)

The rover can autonomously recognize the failure of devices, and in some cases compensate for those failures. The "faster, better, cheaper" approach to spacecraft design has required an acceptance of higher risk, with few fully redundant components. Each time the rover performs an internal health check, it will increment a "failure counter" for each apparently failed device; once the failure count is high enough, no attempt will be made to rely on that device during command execution. If the device begins operating again, the failure count will decrement. If a device draws excessive current when powered on, it will be immediately shutdown and marked as unusable; based on later test results, this status can be cleared by operator commands.

Some fall-back approaches to device failures are used. If the turn rate sensor fails, turn angles are estimated using the wheel encoders. If potentiometers or encoders cease to operate, time-based steering or driving is used.

Rover resource management is similar to that performed by other spacecraft. The rover monitors its internal temperatures on a regular basis between command executions, turning its heaters on or off as necessary. The intent of the thermal control function is to maintain the WEB temperature between -40°C and +40°C throughout the rover mission. This is accomplished by heating the WEB to nearly 40°C by the end of the sol's daytime activities; once the rover shuts down, the WEB will cool through the Martian night, bringing its internal temperature down to approximately -25° C by the time the rover resumes operation the next morning.

The rover also determines the solar power available at a given time, and assesses whether there is sufficient power to execute the next command. By command, battery usage for certain devices can be enabled.

Rover Navigation

The "Go to Waypoint" command is the primary implementation of autonomous rover navigation. The rover operates in a coordinate frame (the "surface-fixed frame") that became fixed to the surface of Mars at the time the lander completed sun-finding and identified the direction of Martian north on sol 1. (The origin of the frame is nominally at the center of the base of the lander.) The human operator specifies the x,y coordinate of a site of interest (e.g., a rock APXS target) in this frame; in addition, the operator specifies the maximum time the rover may take to execute the traverse to this location before the command times out. Intermediate waypoints are also defined if there is a preferential path toward the final destination (e.g., obvious hazards to be avoided, or desired imaging locations along the path). If the rover is not already facing the next waypoint, it will drive in an arcing turn toward the goal, until it is facing the destination. It will then drive an approximately straight line, adjusting its path when it detects drift off its course or encounters a hazard condition.

The rover can identify several types of hazards. They include proximity-detected rocks, drop-offs, and slopes; excessive tilt of the vehicle; triggered contact sensor; loss of communications; motor stall; and the lander itself as a prestored "virtual hazard." Detection of all of these potential hazards is enabled during execution of the "Go to Waypoint" command; however, the specific hazards that the rover is allowed to avoid autonomously (without aborting the traverse) are specified by a settable parameter. (Most conservatively, only autonomous avoidance of proximity hazards is enabled. In rough terrain, contact sensor recovery is also commonly enabled. In contingency sequences, we instruct the rover to avoid all hazard types autonomously as necessary.) If the rover detects a proximity hazard, the vehicle turns in place in increments, until the hazard is no longer detectable. Then the vehicle drives forward one-half vehicle length, after which it resumes normal traverse operations, heading back towards the goal location. At this point, the rover has no memory of the hazard that it has just avoided; it does not maintain a permanent map of the terrain through which it traverses. (Reliance on a map based on outdated information as dead reckoning error accumulates could prove more hazardous to the rover than helpful.)

The average rover traverse in the first 60 sols has been about 2 to 3 meters per sol. (Longer traverses, on the order of 10 meters per sol, are planned as technology experiments later in the mission.) The success of a traverse is dependent on the daily update of the rover's position and orientation using the end-of-day lander images, as well as accurate designation of desired destinations by the Earth-based operators. The rover cannot reach a desired destination unless it is provided with both an accurate indication of where it is (including where it is pointing) and where to find its destination.

Proximity hazard detection is performed using the forward cameras and five laser stripers. Every seven centimeters of traverse, the rover stops and executes a sensing cycle. The rover captures an image both with and without a laser active. Selected scanlines from each image are differenced to locate the laser spot in the scene (i.e. the point at which the laser stripe crosses the scanline). (Figure 3 shows the infrared laser stripe as seen by the rover during surface operations.) If the terrain is flat and level, the laser spot will be visible in a known position along the scanline. Deviations from flat and level ground will cause the laser spot to slide along the scanline, indicating a rock or depression. If the spot cannot be found in the difference image, a significant drop-off may exist. Repeating this process for 5 lasers and four sets of scanlines per difference image generates a set of 20 terrain height measurements. Height differences between adjacent measurements can indicate a rock or hole; sufficient height difference between the lowest and highest measurements in the set indicates a steep slope. False hazard detections can occur if the camera view of a laser spot is blocked by a craggy surface, so ignoring small numbers of data drop-outs is possible by modifying parameter settings in appropriate terrains. During outdoor testing, as well as operations on Mars, the rover has commonly been directed to accept up to three data drop-outs before initiating hazardavoidancebehavior.

Figure 3. Image of laser stripe from front left rover camera

The geometry of the laser stripes has been arranged so that obstacles can be detected to the sides of the rover traverse direction at sufficient range to validate that the entire rover's turning circle is free of hazards. This means that the rover will nominally maintain enough free space around itself to allow for avoiding obstacles detected ahead of it by turning in place and driving forward. This avoids the necessity to drive backwards, since the rover has no proximity hazard detection to the rear. If the density of hazards in the terrain is too high to permit the vehicle to maintain a clear turning circle, a "thread the needle" approach can be enabled by parameter setting adjustment. This technique permits the rover to drive between obstacles that are just further apart than the vehicle width. The rover attempts to drive in a straight line along the perpendicular bisector between the two obstacles, and if it finds a clearing large enough to turn around in before a specified elapsed distance, then it continues on. Otherwise, is backs straight out to the point at which the "thread the needle" behavior was triggered, and then tries another direction.

We have implemented several navigation safety features to protect the rover during waypoint traverses:

To ensure that the rover does not inadvertently traverse beyond communications range, it stops periodically (about once per vehicle length of traverse) to perform a "heartbeat" communications test. If the lander responds, the rover resumes its traverse. Otherwise, the rover retreats 30 cm, turns 45 degrees, and attempts to reestablish contact with the lander.

The lander itself is a potentially serious hazard to the rover. Cleats on the rover's wheels could catch airbag material, possibly permanently entangling the vehicle. A settable parameter permits the human operators to specify just how close to the lander the rover is allowed to go. This virtual hazard is triggered only if the rover is within the hazard radius and driving towards the lander. If the rover is inside the danger zone, but driving away from the lander, it will perceive no hazard. Again, depending on parameter settings, the rover will either autonomously avoid the lander, or abort the remaining traverse.

Contact sensors are located on bumpers on the front and rear of the rover solar panel, and on the lower front body of the rover. Additional contact sensors are incorporated into the APXS deployment mechanism, which is located at the rear of the rover. If an obstacle in the rover's path is not detected by the proximity hazard detection system, triggering any of the bumper contact sensors will either abort the traverse or cause the rover to back up, turn, and avoid the hazard.

If a specified waypoint destination is not reached within the time allotted in the command, the command will time out, setting an error flag. The time limit on command execution prevents the rover from continuing unproductive attempts to achieve an unreachable goal. Depending on the parameter settings in the sequence, any remaining traverse commands will be skipped (since the rover is not where it was expected to be), or the rover will continue on to the next specified location, which may be reachable.

The "Find Rock" command allows the rover to zero in on a rock target at the end of a traverse, autonomously correcting for possible dead reckoning error. The usual strategy to reach a specific rock for the collection of APXS data is to first traverse to the vicinity of the target via one or more "Go to Waypoint" commands. Once there, the rover is commanded to execute a "Turn Toward" the expected rock location, so that the rover is now facing in the direction the rock is most likely to be found. (The "Go to Waypoint" command does not specify the final heading of the rover at the end of a traverse.) The rover then executes a "Find Rock" command specifying coordinates beyond the rock's actual position. The "Find Rock" executes in the same way as a "Go to Waypoint," except that the first time a rock hazard is found during its traverse, the rover will stop, then turn in place while using its hazard detection sensors to determine the extent of the object, and finally turn to face the center of the rock.

To provide full flexibility for rover control, low level motion commands are available. Table 1 lists the full rover command set. (Commands marked with a * are vehicle motion commands.) The "Move" command directs the rover to drive forward or backward with fixed wheel steering angles for a specified number of centimeters. Variations of the "Turn" command allow operators to specify relative turns, turns to absolute heading, and turns to face a particular coordinate location in the "surface-fixed frame." Additionally, selected hazard detection capabilities can be disabled during "Go to Waypoint" commands if a specific circumstance so indicates.

Contingency Scenario Responses

Earth to Lander Command Loss--The Mars Pathfinder mission planners prepared for the possibility of loss of communications with the spacecraft immediately after landing. If a two-way loss of communications had occurred, then no surface mission would have been possible (or any results of such a mission would never be known). However, if only the receiver on the lander had failed, then telemetry from the spacecraft would still be received on the ground; only the opportunity to command the spacecraft (and rover) would have been lost. In order to perform a useful mission under such a "command loss" scenario, a Backup Mission Load (BML) was designed and uplinked to the spacecraft during the cruise mission phase. The BML includes a set of command sequences for both the lander and the rover. The BML would have been activated after sufficient time (2 sols) had elapsed since the lander last received any sequences from the Earth. The lander would then have released sequences to the rover to stand up, deploy down the lander ramp, and perform surface operations. The lander would have transmitted telemetry, with the hope that it will be received by the DSN.

Since, in this scenario, the rover would still regularly receive new sequences, it would continue to operate in a nominal mode. The BML would have allowed some coordination of lander and rover activities during the early part of the mission. For the first few sols, until dead reckoning error accumulated to significant levels, the lander should have been able to point its camera to image the rover at its end-of-day location. When rover sequences in the BML would no longer be useful (i.e., when the rover's location was effectively unknown), the rover would have been allowed to trigger its own onboard contingency sequence.

Backup Mission Load sequences were loaded onboard the Pathfinder lander to support command-loss scenarios triggered at any time between landing and sol 60. With the successful completion of both the lander and rover primary missions, there is no longer significant utility in a generic mission performed via a set of "canned" command sequences.

Lander to Rover Command Loss--Neither the BML or nominal rover command sequences would have been effective means of commanding the rover if the UHF communications link between the lander and rover had failed on or after landing. In response to this possible scenario, the rover team developed a Contingency Mission Load (CML) which was placed into the rover's non-volatile memory before delivery of the rover to Kennedy Space Center for launch preparation.

If the failure of the lander/rover link had been bi-directional, then the only telemetry documenting rover activities would have been lander imaging of rover traverses. However, if only the command link between the rover and lander had failed, downlink of rover telemetry would still have been possible, although coordination of activities between the two spacecraft would no longer be feasible. The rover software was designed to activate a contingency sequence if the rover fails to receive a complete command sequence for approximately two sols. The particular sequence to be triggered depends on the mission phase of the rover (i.e., prelaunch, cruise, prerelease, predeploy, primary, or extended). For example, if the rover were still on the lander petal when the contingency mission activated, the triggered sequence would cause it to stand up, then switch to the next phase for driving down the ramp. If the rover were already performing surface operations when the communications link failed, the activated sequence would continue surface operations, attempting to circumnavigate the lander at a range of approximately 5 meters, while finding rocks, taking APXS measurements, performing MAE, WAE, and soil experiments, and imaging. Each telemetry frame generated by the sequence would be transmitted twice. Without handshaking, the lander has no mechanism to

Table 1. Rover Command Set
(Commands marked with a * are vehicle motion commands.)

ABORT SEQUENCE
APPLY PATCH SET s LENGTH n
CALL FUNCTION AT ADDRESS a WITH BANKS b1, b2 [PARAMETERS p1 p2 ...] (RL)
CAPTURE IMAGE WITH CAMERA c AT EXPOSURE t, RETURN REGION FROM (r1,c1) TO (r2,c2) WITH APID a[compressed]
CLEAR c
DRIVE MOTORS*
DEPLOY APXS p
END OF SEQUENCE
FIND ROCK NEAR x y WITHIN m MINUTES*
GO TO WAYPOINT AT x y WITHIN m MINUTES*
HEALTH CHECK AT LEVEL n
HEAT FOR n MINUTES IN SEQUENCE s [s ...]
LIMIT-CALIBRATE POSITION SENSORS
MATERIAL ADHERENCE EXPERIMENT
MOVE BACKWARD n COUNTS*
MOVE FORWARD n COUNTS*
OUTPUT x TO PORT p
PATCH MEMORY SET s CRC c SEGMENT b ADDRESS a TO x x x ...
POSITION APXS*
READ ANALOG CHANNEL c AT GAIN g [c AT g ...]
READ BYTES FROM PORT p p p ...
READ n BYTES FROM SEGMENT s ADDRESS a
RUN MOTOR m FOR t CENTONS (Fractional Seconds)
RUN MOTOR m TO p (Counts)
SEND APXS COMMAND c
SET CLOCK TO t
SET DEVICE d STATUS s
SET ERROR MASK TO m
SET MISSION PHASE p
SET PARAMETER p = value
SET VEHICLE POSITION TO x y, HEADING TO h
SHUTDOWN UNTIL t
SOIL MECHANICS TEST ON WHEEL w AT POSITION p, RUN FOR n COUNTS, REPEAT k TIMES
SYNCHRONIZE CLOCK
TEST n BYTES OF MEMORY IN SEGMENT s AT ADDRESS a
TURN LEFT n BAMS (Binary Angular Measurements)*
TURN OFF DEVICES d
TURN ON DEVICES d
TURN RIGHT n BAMS (Binary Angular Measurements)*
TURN TO HEADING h*
TURN TOWARD x, y*
UNSTOW p d c l r *
WAIT FOR s SECONDS
WAIT UNTIL LOCAL TIME t
WAIT UNTIL MISSION TIME t
WAIT UNTIL sensor_value < value, LIMIT m MINUTES
WAIT UNTIL sensor_value > value, LIMIT m MINUTES
WHEEL ABRASION MEASUREMENT AT GAIN g
ZERO-CALIBRATE POSITION SENSORS

determine into what kind of rover packet to reassemble the telemetry, so all telemetry would be classified as "unrecognized rover packets" and forwarded to the ground for reconstruction. If communications were reestablished, the rover would resume normal operation immediately when a command sequence was successfully received from the lander.

As has been the case with the Backup Mission Load, the need for the Contingency Mission Load has decreased with the continued success of the rover's mission on Mars. On sol 80, the contingency command sequence for the "primary" mission phase was replaced with a new load, which is essentially a multi-sol safing sequence. If the rover fails to receive a new command sequence after several days in the contingency mode, it will then activate the contingency sequence associated with the "extended" mission phase, which will cause it to attempt to drive toward the lander, under the assumption that the rover may have driven out of communications range or into a communications null region.

On sol 83, contact with the Pathfinder lander was lost. While attempts to regain communications with the spacecraft continue, we can only presume that the rover has properly activated its contingency sequences. If so, it began driving towards the lander on sol 92. Since the lander is a virtual hazard representing a keep-out zone, the rover should continue to circle the lander at a range of approximately three meters until communications is reestablished or a hardware failure occurs.

7. PERFORMANCE ON MARS

The overall performance of Sojourner on Mars has exceeded both design goals and expectations, as evidenced by the successful accomplishment of all rover mission objectives and the continuation of the rover's activities beyond even the planned extended mission duration. As of sol 83, 16 distinct sites (9 rocks, 7 soil locations) had been analyzed by the APXS; over 120 MAE and 9 WAE experiments had been performed; and more than 500 images had been captured. The rover had traversed over 100 meters total integrated distance, nearly circumnavigating the lander (Figure 4) and averaging 2.7 meters per traverse day (Figure 5). (Several sols included no traverses, often because of either scheduled or unscheduled loss of uplink opportunities between Earth and the lander. In these cases, the rover either continued ongoing APXS data collection, or executed a runout sequence leaving the rover in a safe state while waiting for the next sequence to arrive.) The rover's traverse capabilities have allowed us to both view features at higher resolution than would be possible from the lander's camera, and to image features that cannot be seen at all from the lander's vantage point (see Figures 6 and 7).

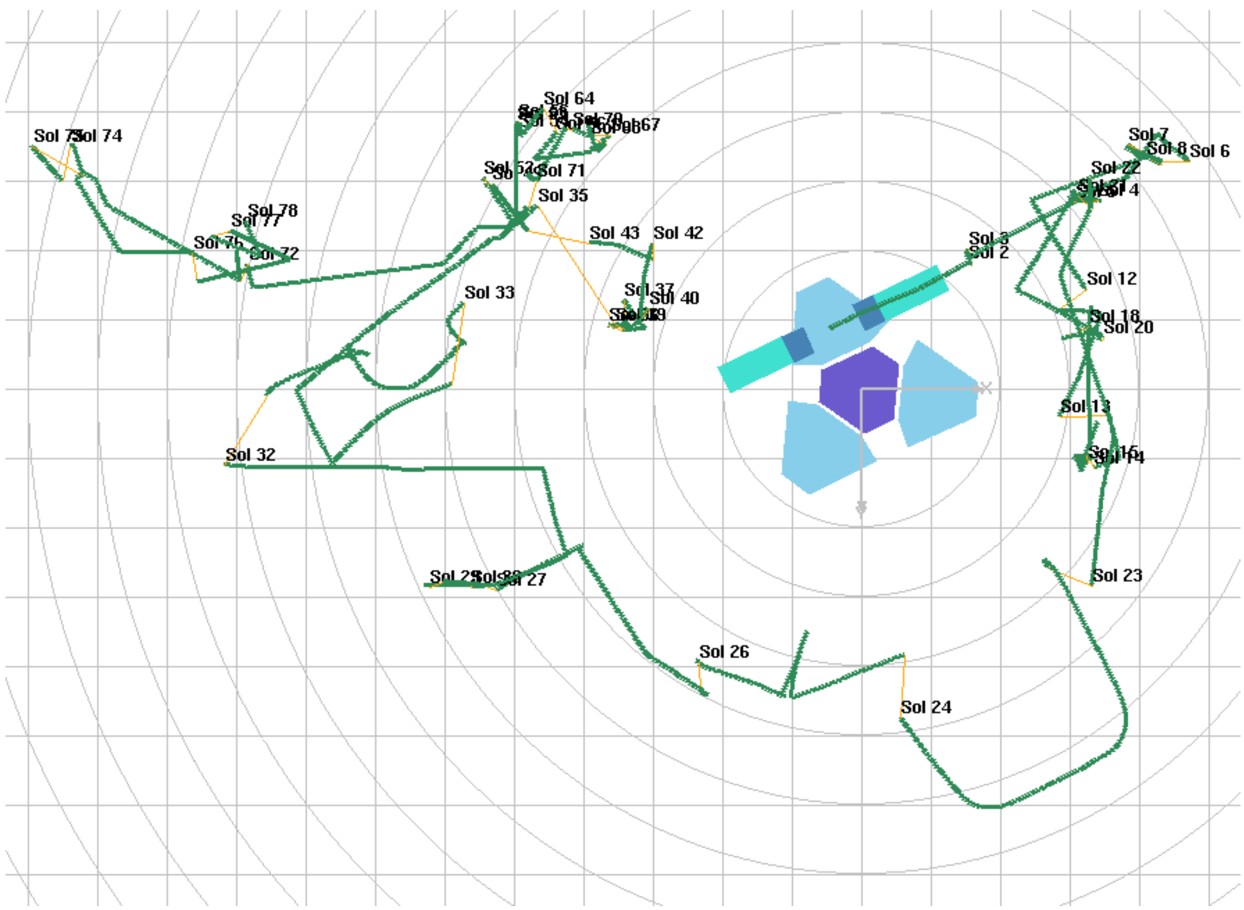

Figure 4. Rover traverses during first 78 sols
(grid spacing is 1 meter)

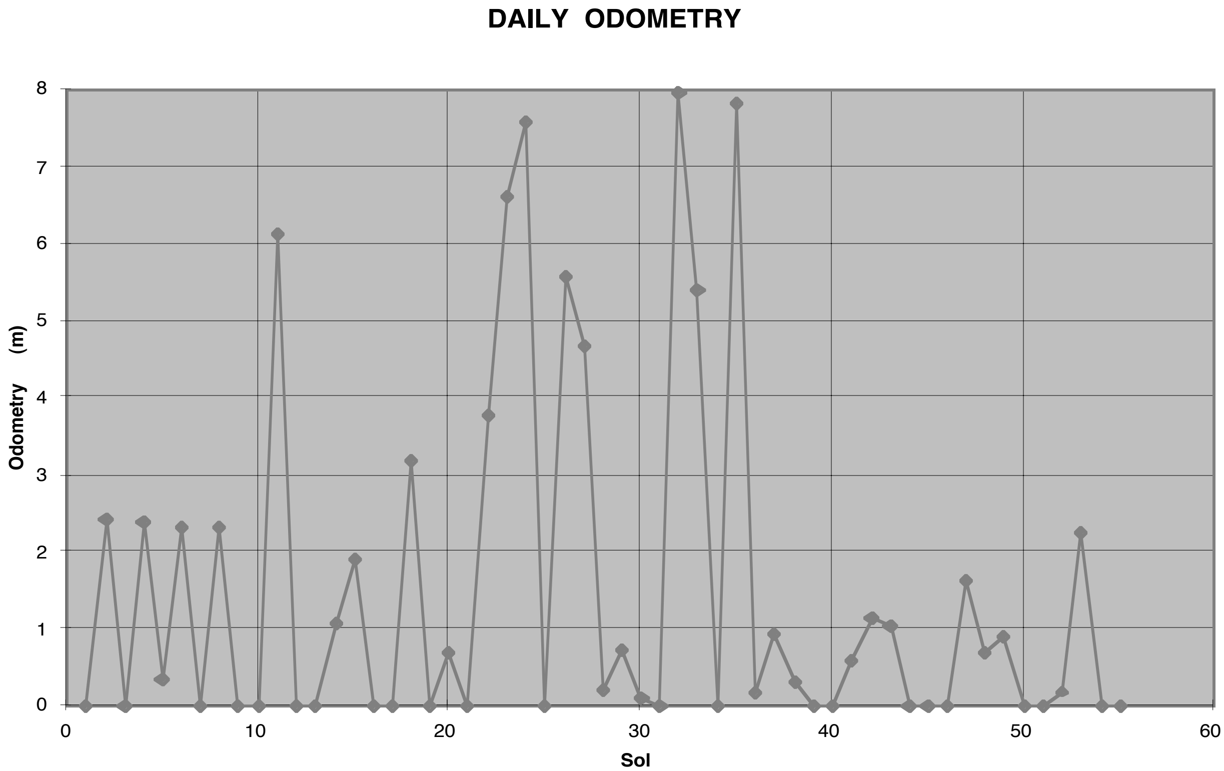

Figure 5. Rover Odometry per Sol

Figure 6. Closeup rover image of the rock “Chimp”

Figure 7. Sand dunes behind the “Rock Garden” visible only from the rover

Figure 8. Initial rover traverse and smooth terrain near lander

The autonomous navigation performance of the rover on Mars has generally been equal to or better than the performance observed using the testbed rover during Operations Readiness Tests on Earth.

Because of the nearly obstacle-free nature of the terrain in the immediate vicinity of the lander ramp (see Figure 8), initial rover traverses were commanded as low-level moves, with no “Go to Waypoint” commands used. We also wished to avoid reliance on waypoint traverses until we had evaluated both the rover's dead reckoning performance in the Martian terrain, and the ability of the laser/camera hazard detection subsystem to detect the laser stripes under Mars illumination and albedo conditions.

Rover camera images captured with hazard detection lasers powered on (e.g., Figure 3) confirmed the visibility of the stripes. The first "Go to Waypoint" command was executed on sol 12.

Our overall experience with rover navigation was that while the dead reckoning performance was poor, hazard detection and avoidance worked well. Consistent with earlier ground testing, position error was roughly 5-10% of distance traveled, and average drift of the heading reference subsystem was approximately 13 degrees/sol of traverse. The result of this dead reckoning performance was that when autonomous traverse was enabled, the "Go to Waypoint" commands did not always lead the rover to the expected location, but the rover nevertheless successfully avoided non-traversable hazards. In one instance, the rover "threaded the needle" between two hazards which were barely more than a rover width apart. This led the rover into a region near the rock Wedge (Figure 9). Subsequent attempts to use low-level "Move" and "Turn" commands to exit the region required several sols to finally move away from Wedge and into the vicinity of the aptly-named "Rock Garden." (Low-level commands near Wedge often resulted in abbreviated traverses because of tilt or articulation hazard conditions triggered when the rover rode up onto nearby rocks. Misregistration of as little as 10 centimeters of the stereo-derived terrain database and the end-of-day images of the rover contributed to the problem of encountering rather than avoiding obstacles.)

Further complicating traverse operations was the occasional spurious measurements obtained from of two out of three onboard accelerometers used to determine the vehicle's orientation. On some sols these accelerometers would intermittently generate spurious values in error by tens of degrees. (This behavior of the accelerometers had not been observed prior to landing, either on the flight rover or the ground test unit.) While the rover's sensor polling software routinely filtered out single false values, these erroneous readings sometimes persisted long enough to trigger a tilt hazard and bring the sol's traverse to a halt. The accelerometers could be disabled to prevent unwarranted responses to nonexistent hazards; but this meant that the rover would be unlikely to recognize a true tilt threat if it arose. In some cases, we have needed to disable the accelerometers under exactly the terrain conditions when they would be most useful. Selecting the appropriate state for the accelerometers became part of the planning process of each rover traverse.

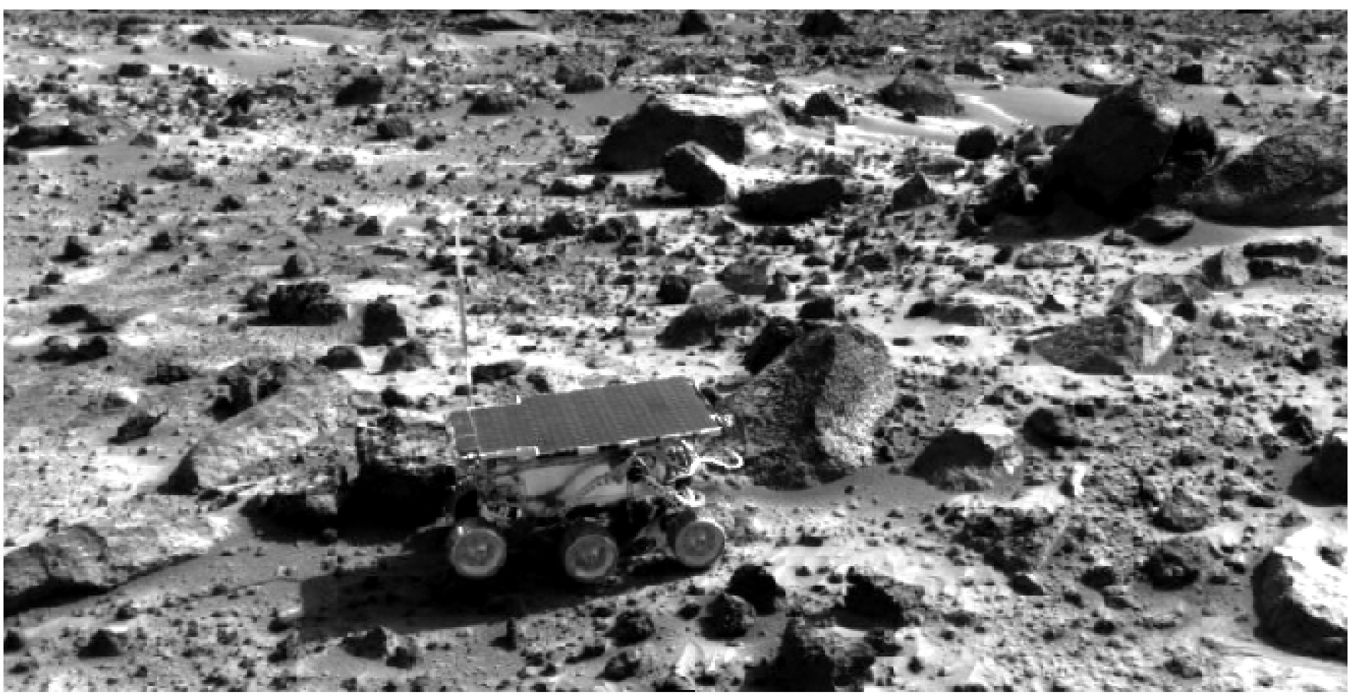

Figure 9. Sojourner to the left of the rock "Wedge" on Sol 35.
Part of the "Rock Garden" is visible in the upper right of the image.

(a)

(b)

(c)

(d)

Figure 10. Rover traverse over rock on Sol 24

Despite some difficulties in operating the rover, Sojourner demonstrated its ability to drive in the terrain of the landing site. Rover "movies" were regularly constructed by capturing a series of images using the lander camera pointed where the rover was expected to drive at a particular time of day. An excerpt from one of these movies is shown in Figure 10, clearly indicating the rover's capability to negotiate a rock nearly a wheel diameter in height. In this series of images the rear of the rover is clearly visible, with the APXS instrument fully retracted; the rover is driving away from the camera.

8. CONCLUSIONS

Unlike spacecraft developed for previous planetary exploration missions, Sojourner operates in a non-deterministic environment, in which each step may yield unexpected results because of unknown terrain conditions. Although modest in capability and complexity, the microrover is unique among robotic missions to date in its ability to operate in an unmodeled environment and choose actions based on sensor input to accomplish requested objectives. As such, Sojourner is probably the most autonomous deep space probe yet launched.

The autonomous navigation capabilities of the Sojourner rover have proven sufficient to reach the sites of interest at the Pathfinder landing site. Traverses in the smoother areas within the site have been straightforward; navigation through the rockier areas, most notably the region dubbed the "Rock Garden," has been more problematic, requiring several sols to cover a few meters of obstacle-strewn terrain. While some of the observed difficulties are clearly due to limitations in the implementation of autonomous navigation onboard the vehicle, much can be attributed to the caution of the rover team in enabling the rover's full suite of hazard avoidance features during specific traverses. This caution is understandable, given that each rover traverse inherently puts the vehicle at risk, and the consequence of a poorly commanded traverse may be the premature end of the mission.

Future planned rover missions, such as the Mars Surveyor Program 2001 mission, will not have the luxury of accomplishing their objectives while maintaining such a conservative approach to risk. In these missions, the rover will be required to traverse approximately 100 meters per sol in order to reach sites of scientific interest and collect samples for eventual return to Earth. Such rover navigation performance is equivalent to performing all of the traverses of the Sojourner rover during the entire Pathfinder surface mission to date in a single sol. While such long distance traverses will clearly entail a significant increase in required autonomous capability for future rovers, Sojourner has already proven the feasibility and value of mobile robots for planetary surface exploration.

ACKNOWLEDGMENTS

The work described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration.

The Microrover Flight Experiment (MFEX) is a NASA OACT (Office of Advanced Concepts and Technology) activity.

We would like to acknowledge the efforts of the engineers and scientists of the Mars Pathfinder team, who together made the mission a success beyond our expectations.

In addition, we would like to thank the members of the IMP (Imager for Mars Pathfinder) team for the IMP images of the rover and the surface of Mars (figures 8, 9, and 10).

Figure 4 was provided by Richard Welch of the Pathfinder rover team.

REFERENCES

[1] Henry W. Stone, "Mars Pathfinder Microrover: A Low-Cost, Low-Power Spacecraft," Proceedings of the 1996 AIAA Forum on Advanced Developments in Space Robotics, Madison, WI, August 1996.

[2] Jake Matijevic, "Mars Pathfinder Microrover - Implementing a Low Cost Planetary Mission Experiment," Proceedings of the Second IAA International Conference on Low-Cost Planetary Missions, John Hopkins University Applied Physics Laboratory, Laurel, Maryland, April 1996.

[3] Jack Morrison, and Tam Nguyen, "On-Board Software for the Mars Pathfinder Microrover," Proceedings of the Second IAA International Conference on Low-Cost Planetary Missions, John Hopkins University Applied Physics Laboratory, Laurel, Maryland, April 1996.

[4] Henry W. Stone, "Design and Control of the MESUR/Pathfinder Microrover," Proceedings of the 1993 International Conference on Advanced Robotics, Tokyo, Japan, November 1993.

Andrew Mishkin *is the Pathfinder Microrover Mission Operations Engineer, responsible for building the rover operations team, interacting with Pathfinder scientists regardingrover planning, and designing rover command sequences . He graduated from UCLA in Los Angeles, California with a BS and MS in systems engineering. He is a Group Leader for the Robotic Vehicles group at the Jet Propulsion Laboratory, where he has coordinated the development and integration of several planetary rover efforts since 1985. He is currently system engineer for the Mars Surveyor Program 2001 rover mission, and was system engineer for the Control and Navigation subsystem for the Sojourner rover. In his spare time he enjoys hiking, writing, and sculpting in bronze, neon, and steel.*

Jack Morrison *was the primary software designer for the Sojourner rover. During Pathfinder mission operations, he has been one of two primary rover drivers designating vehicle traverses. He is a native of Southern California, with a 1978 BS in Math from UCLA. His career has taken him on adventures in aerospace and commercial development, from air defense systems to video effects boxes; from 8-bit microcontrollers to gigabyte* *supercomputers; from air force bases to investment firms, and from Venus to Mars. His articles have been published in computer graphics books and his photo with a Cabbage Patch doll published on the Web. His hobbies include virtual air combat and literal equestrian jumping.*

Tam Nguyen *is a software engineer for the Robotic Vehicles Group at the Jet Propulsion Laboratory. Currently, he is a software engineer for the Microrover of the Mars Pathfinder Project and also a member of the Mars Pathfinder Microrover Downlink Data Analysis team. Since 1988, he has developed software and integrated hardwarein the areasof motion control and navigation for several planetary rover research programs. Prior to these programs, he was a key software engineer for the Three Axis Levitator Project, a physics flight experiment flown on the Shuttle. Nguyen has a BS in Physics and Chemistry and a MSEE from CSU Long Beach, California.*

Henry Stone *receivedhis BSEE, MSEE, and PhD from Carnegie Mellon University, in Pittsburgh, Pennsylvania, in 1981, 1983, and 1986, respectively. His dissertation involved the kinematic modeling, identification, and control of robotics manipulators. He joined the Jet Propulsion Laboratory, Pasadena, Ca, in 1986 and initially worked on the Telerobotics Testbed Project which was aimed at developing technologies for controlling multiple manipulators for in-flight servicing of Earth-orbiting spacecraft/satellites. In 1991 he began working in the area of mobile robots and became the Cognizant Engineer of the Hazbot Vehicle, a terrestrial robot designed to remotely investigate and cleanup hazardous material spills. In 1993 he became the Technical Manager of the Mars Pathfinder Microrover's Control and Navigation Subsystem. He has recently been selected to lead the design of the Control and Navigation Subsystem for the 2001 Rover (i.e., the long range rover which will fly to Mars aboard NASA's 2001 Mars Mission).*

Brian Cooper *was the principal developer of the Rover Control Workstation for the MFEX rover, as well as the primary rover driver during surface operations. He graduatedfrom the University of California, Irvine with a BS in ElectricalEngineering in 1981, and earned an MS in Computer Engineering from the University of Southern California in 1985. He served as an officer in the United States Air Force from 1981-1985, stationed at Los Angeles Air Force Station (now LA Air Force Base), El Segundo, CA where he worked on military satellite systems. He joined JPL in 1985 as an engineer working on robotic vehicle research,in the areaof computer-aided remote driving. He developed the user interface and control station for numerous robotic vehicles, leading up to this assignment on Mars Pathfinder. Brian's major hobbies include flying, scuba diving, magic (he's been a member of the Magic Castle in Hollywood, CA for over 20 years), and (most important) building his own airplane.*

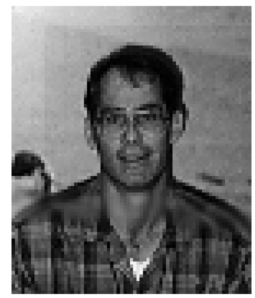

Brian Wilcox *is the Supervisor of the Robotic Vehicles Group at JPL. He originated the laser striping hazard avoidance and imaging sensor used on Sojourner, as well as the basic hazard avoidancealgorithm. He has been developing planetary rovers at JPL since 1982 and has been group supervisor since 1985. He was awarded the NASA Exceptional Engineering Achievement medal in 1992 for his contributions to planetary rover research. He currently also manages the development of the nanorover (a rover with a mass of a few hundredgrams) for the MUSES-CN asteroid mission scheduled to launch in early 2002. He has a B.S. in Physics and a B.A. in Mathematics from the University of California at Santa Barbara and an M.S. in Electrical Engineering from the University of Southern California.*